- We have tailored resource policies to aid users whenever possible within our research mandate and available facilities. Our approach to system scheduling, overload control, file space management, etc. all attempt to give users the greatest latitude possible to pursue their research goals consistent with fairly meeting our responsibilities in administering SUMEX as a national resource.

III.D.2. 2060 Cost Center

General Cost Center Structure

Our renewal proposal for the five-year period 8/1/86-7/31/91, submitted to the Division of Research Resources in June 1985, called for phasing out NIH support for DEC 2060 mainframe operations over the course of the grant period and the establishment of a cost center at Stanford to recover the unsubsidized costs of 2060 operations from the user community. This phasing-out process is taking place linearly over five years, with 20% of the 2060 costs being recovered in renewal year 1 (Grant Year 14), 40% in year 2, 60% in year 3, 80% in year 4, and 100% starting in year 5. In this process, we are attempting to minimize the barriers for national projects by using the continuing partial BRTP subsidy to cover their costs for as long as possible. In this past year, use of the 2060 by members of the national AIM community has been free of charge. Thus, the Stanford user projects are bearing the entire brunt of cost recovery during the first few years. Our plan is conservative, however, in that we are doing this gradually and responsibly so that our users can secure the funding resources and make software changes necessary to allow them to relocate to other facilities or move to workstation environments for their research.

To implement this plan, during the summer of 1986, we requested and received approval from the Government Cost and Rate Studies section of Stanford's Controller's Office to establish a 2060 cost center effective August 1, 1986. We set up the cost center with the simplest possible charge structure in order to minimize the accounting and administrative overhead, establishing a charge rate per CPU hour based on our projections of 2060 operations costs and anticipated billable Stanford project CPU usage. The initial rate was established at $95 per CPU hour.

We closely monitored the cost center expenses and revenues during the year. A mid-year analysis of cost center performance indicated that expenses would be somewhat lower and billable CPU usage somewhat higher than originally projected. To produce a year-end (July 31, 1987) break-even condition for the cost center, we lowered the charge rate as of February 1 to $75 per CPU hour. Figure 17 shows the cumulative user revenues collected by month for the period August 1986 through April 1987 as well as the ideal (linear) cost center recovery line.

The cost center rate for Stanford users is expected to increase substantially at the beginning of each succeeding grant year through renewal year 5, as NIH subsidy of 2060 costs is incrementally withdrawn.

Remote Network Costs

Until this year, the costs associated with networking were supported by NIH through Rutgers University. Beginning this grant year, however, NIH is funding our networking costs directly as part of our 2060 operations budget, and we have entered into a contract with TELENET Communications Corporation for networking services. To underscore our commitment to subsidize the national AIM community's 2060 usage as long as possible, we have been paying for TELENET services directly from the SUMEX

grant this year on the assumption that national community members would represent the vast majority of TELENET users. However, all other 2060-related expenses are charged directly to the cost center and then charged out to Stanford users according to their CPU usage and to the SUMEX grant in keeping with its level of subsidy of 2060 operations.

This early practice of paying for TELENET services directly from the grant has complicated our accounting procedures, since networking expenses must ultimately be taken into consideration in allocating total annual 2060 operations costs in correct proportions to the resource budget and to Stanford users. Also, a recent analysis of our networking usage indicated that the use of TELENET by Stanford groups is considerably higher than expected. Therefore, since networking services are not being used exclusively by the national user community as originally believed, we plan to change our procedure and charge TELENET costs directly to the cost center in future years.

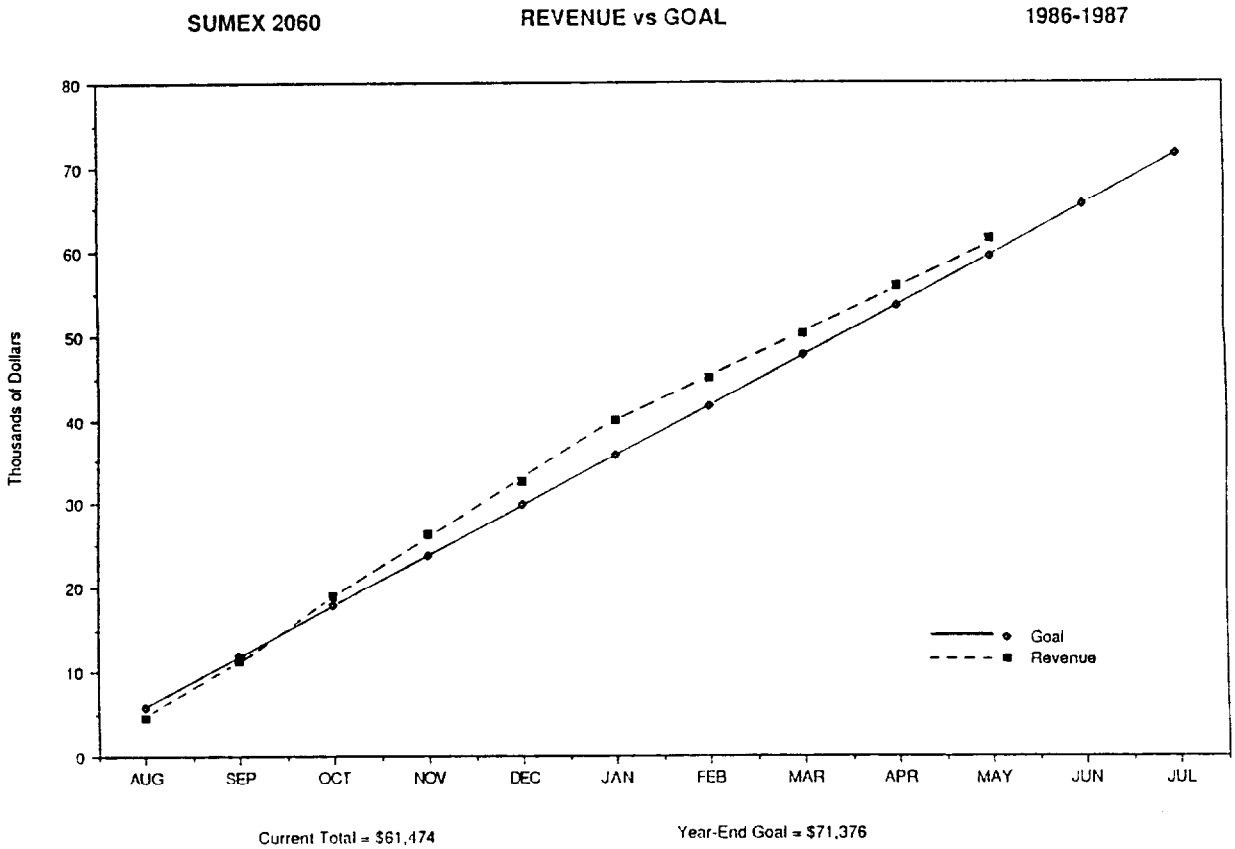

Figure 17: 2060 Cost Center Performance

III.E. Dissemination of Resource Information

We are continuing our past practice of making a substantial effort to disseminate the AI technology developed here. This has taken the form of many publications -- over forty-five combined books and papers are published per year by the KSL; wide distribution of our software including systems software and AI application and tool software, both to other research laboratories and for commercial development; production of films and video tapes depicting aspects of our work; and significant project efforts at studying the dissemination of individual applications systems such as the GENET community (DNA sequence analysis software) and the ONCOCIN resource-related research project (see 123).

Software Distribution

We have widely distributed both our system software and our AI tool software. Since much of our general system-level software is distributed via the ARPANET we do not have complete records of the extent of the distribution. Software such as TOPS-20 monitor enhancements, the Ethernet gateway and TIP programs, the SEAGATE AppleBus to Ethernet gateway, the PUP Leaf server, the SUMACC development system for Macintosh workstations, and our Lisp workstation programs are frequently distributed in this manner to the ARPANET community and beyond. Since our SUMACC software development system for Macintosh workstations is considered to be in the "public domain", we have turned it over to Information Analysis Associates, Mountain View, CA. for distribution (for a minimum charge) to groups not associated with the ARPANET.

Our primary distribution effort is directed towards our AI tool material. In recent years, the volume of inquiries for this type of software and requests for tapes has been a substantial burden on the staff. Records indicate that over the past three years there have been about 1,050 inquiries that have resulted in the distribution of written material about our software systems. It is likely that there have been a similar number of unrecorded or informal interactions on the part of the staff. It was therefore decided to turn over most of this type of software distribution to Stanford's Office of Technology Licensing (OTL).

This organization handles software distribution and technology licensing matters for much of the Stanford community. Since there are several OTL staff members assigned to the distribution of Stanford software, requests for information and tapes are handled quickly and efficiently. Also, OTL's staff has the expertise needed to handle the legal questions that frequently arise in the distribution of software, and an established computerized record-keeping scheme. SUMEX staff continues to be available as needed to assist OTL with special administrative and technical matters.

Unfortunately, start-up delays in the transfer of software distribution to the Office of Technology Licensing and the preparation of new versions of MRS and BB1 have temporarily reduced our distribution volume. During this report period we distributed eleven copies of MRS, eight copies of BB1 and one each of AGE, EMYCIN, GENOA, and CONGEN. During the past year the reconstruction of the distribution packages for the DENDRAL Project (GENOA and CONGEN) has been completed. In December of this year, a five-year exclusive licensing agreement (with Molecular Designs, Ltd.) for the DENDRAL material will expire, and we will therefore have more flexibility in distributing this material.

We continue to make a special effort to assist other members of the SUMEX-AIM community in integrating the technologies needed for biomedical AI research. This is often achieved through direct contact with staff members at these institutions at meetings and workshops or via electronic mailing lists. For example, the Info-1100

mailing list, which is maintained at SUMEX-AIM, has several hundred members (users of Xerox 1100 Series equipment) and is monitored by our staff. This list is used to distribute things like hardware and software bug reports and fixes and system tools and is very valuable to the AIM community Interlisp users.

Video Tapes and Films

The KSL and the ONCOCIN project have prepared several video tapes that provide an overview of the research and research methodologies underlying our work and that demonstrate the capabilities of particular systems. These tapes are available through our groups, the Fleischmann Learning Center at the Stanford Medical Center, and the Stanford Computer Forum, and copies have been mailed to program offices of our various funding sponsors. The three tapes include:

- *Knowledge Engineering in the Heuristic Programming Project* -- This 20-minute film/tape illustrates key ideas in knowledge-based system design and implementation, using examples from ONCOCIN, PROTEAN, and knowledge-based VLSI design systems. It describes the research environment of the KSL and lays out the methodologies of our work and the long-term research goals that guide it.

- *ONCOCIN Overview* -- This is a 30-minute tape providing an overview of the ONCOCIN project. It gives an historical context for the work, discusses the clinical problem and the setting in which the prototype system is being used, and outlines the plans for transferring the system to run on single-user workstations. Brief illustrations of the graphics capabilities of ONCOCIN on a Lisp workstation are also provided.

- *ONCOCIN Demonstration* -- This 1-hour tape provides detailed examples of the key components of the ONCOCIN system. It begins with a demonstration of the prototype system's performance on a time-shared mainframe computer and then shows each of the elements involved in transferring the system to Lisp workstations.

III.F. Suggestions and Comments

Resource Organization

We continue to believe that the Biomedical Research Technology Program is one of the most effective vehicles for developing and disseminating technological tools for biomedical research. The goals and methods of the program are well-designed to encourage building of the necessary multi-disciplinary groups and merging of the appropriate technological and medical disciplines.

Electronic Communications

SUMEX-AIM has pioneered in developing more effective methods for facilitating scientific communication. Whereas face-to-face contacts continue to play a key role, in the longer-term we feel that computer-based communications will become increasingly important to the NIH and the distributed resources of the biomedical community. We would like to see the BRTP take a more active role in promoting these tools within the NIH and its grantee community.

IV. Description of Scientific Subprojects

The following subsections report on the AIM community of projects and "pilot" efforts including local and national users of the SUMEX-AIM facility at Stanford. However, those projects admitted to the National AIM community which use the Rutgers-AIM resource as their home base are not explicitly reported here.

In addition to these detailed progress reports, abstracts for each project and its individual users are submitted on a separate Scientific Subproject Form. However, we have included here briefer summary abstracts of the fully-authorized projects in Appendix B on page 221.

The collaborative project reports and comments are the result of a solicitation for contributions sent to each of the project Principal Investigators requesting the following information:

I. SUMMARY OF RESEARCH PROGRAM
- A. Project rationale
- B. Medical relevance and collaboration
- C. Highlights of research progress
  - --Accomplishments this past year
  - --Research in progress
- D. List of relevant publications
- E. Funding support

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE
- A. Medical collaborations and program dissemination via SUMEX
- B. Sharing and interactions with other SUMEX-AIM projects (via computing facilities, workshops, personal contacts, etc.)
- C. Critique of resource management (community facilitation, computer services, communications services, capacity, etc.)

III. RESEARCH PLANS
- A. Project goals and plans
  - --Near-term
  - --Long-range
- B. Justification and requirements for continued SUMEX use
- C. Needs and plans for other computing resources beyond SUMEX-AIM
- D. Recommendations for future community and resource development

We believe that the reports of the individual projects speak for themselves as rationales for participation. In any case, the reports are recorded as submitted and are the responsibility of the indicated project leaders. The only exceptions are the respective lists of relevant publications which have been uniformly formatted for parallel reporting on the Scientific Subproject Form.

IV.A. Stanford Projects

The following group of projects is formally approved for access to the Stanford aliquot of the SUMEX-AIM resource. Their access is based on review by the Stanford Advisory Group and approval by Professor Feigenbaum as Principal Investigator.

In addition to the progress reports presented here, abstracts for each project and its individual users are submitted on a separate Scientific Subproject Form.

IV.A.1. GUIDON/NEOMYCIN Project

GUIDON/NEOMYCIN Project

William J. Clancey, Ph.D.
Department Computer Science
Stanford University

Bruce G. Buchanan, Ph.D.
Computer Science Department
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The GUIDON/NEOMYCIN Project is a research program devoted to the development of a knowledge-based tutoring system for application to medicine. The key issue for the GUIDON/NEOMYCIN project is to develop a program that can provide advice similar in quality to that given by human experts, modeling how they structure their knowledge as well as their problem-solving procedures. The consultation program using this knowledge is called NEOMYCIN. NEOMYCIN's knowledge base, designed for use in a teaching application, is the subject material used by a family of instructional programs referred to collectively as GUIDON2. The problem-solving procedures are developed by running test cases through NEOMYCIN and comparing them to expert behavior. Also, we use NEOMYCIN as a test bed for the explanation capabilities incorporated in our instructional programs.

The purpose of the current contracts is to construct a knowledge-based tutoring system that teaches diagnostic strategies explicitly. By strategy, we mean plans for establishing a set of possible diagnoses, focusing on and confirming individual diagnoses, gathering data, and processing new data. The tutorial program has capabilities to recognize these plans, as well as to articulate strategies in explanations about how to do diagnosis. The strategies represented in the program, modeling techniques, and explanation techniques are wholly separate from the knowledge base, so that they can be used with many medical (and non-medical) domains. That is, the target program will be able to be tested with other knowledge bases, using system-building tools that we provide.

B. Medical Relevance and Collaboration

There is a growing realization that medical knowledge, originally codified for the purpose of computer-based consultations, may be used in additional ways that are medically relevant. Using the knowledge to teach medical students is perhaps foremost among these, and GUIDON2 focuses on methods for augmenting clinical knowledge in order to facilitate its use in a tutorial setting. A particularly important aspect of this work is the insight that has been gained regarding the need to structure knowledge differently, and in more detail, when it is being used for different purposes (e.g., teaching as opposed to clinical decision making). It was this aspect of the GUIDON research that led to the development of NEOMYCIN, which is an evolving computational model of medical diagnostic reasoning that we hope will enable us to better understand and teach diagnosis to students. An important additional realization is that these structuring methods are beneficial for improving the problem-solving performance of consultation programs, providing more detailed and abstract explanations to consultation users, and making knowledge bases easier to maintain.

As we move from technological development of explanation and student modeling capabilities, we are now collaborating closely with medical students and physicians to design an effective, useful tutoring program. In particular, medical students have served as research assistants, and a current MSAI student is an experienced physician, John Sotos, from Johns Hopkins. The project also collaborates with a community of researchers focusing on medical education, funded by the Josiah Macy, Jr. Foundation.

C. Highlights of Research Progress

C.1 Accomplishments This Past Year

C.1.1 The GUIDON-DEBUG Tutoring Program

We began 1986 with a concerted effort to construct a tutorial program called GUIDON-DEBUG. The idea behind this system is to have a student debug a faulty knowledge base by using graphic explanation and editing tools. A prototype was demonstrated at the annual ONR conference in March. However, after trials with medical students we realized that 1) it was difficult to choose a fault at the right level of difficulty for a student, and 2) the program lacked ability to help the students and evaluate their debugging because it lacked an internal model of how to debug. We concluded that GUIDON-DEBUG development should be deferred until the proposed knowledge acquisition module (see below) is completed.

C.1.2 The GUIDON-MANAGE Tutoring Program

At this point we returned to an alternative conception described in our original proposal, a program called GUIDON-MANAGE. This program teaches a student the language of diagnosis by having him or her enter all requests for patient information as an *abstraction.* Thus, the student issues "strategic commands" such as "test the hypothesis meningitis" or "ask a follow-up question about the headache," and the program (NEOMYCIN) carries out the tactics. By year end, this program was well along, with a complex interpreter for simulating NEOMYCIN to generate help, a feedback window to indicate what NEOMYCIN did when it carried out the commands, and many menus for making input to the program convenient. Research continues to focus on the assistance and feedback components of the program.

GUIDON-MANAGE is now conceived to be the first step in a three-step tutorial program which will include GUIDON-WATCH (which we previously developed) and a yet to be named tutorial module. In these three steps, the student will solve a problem, watch NEOMYCIN solve a problem, and then explain his solution and seek explanations about NEOMYCIN's solution. In this way, we use the program as a *model* that the student can study and compare to his own reasoning.

C.1.3 The GUIDON-MANAGE Tutoring Program

Research in explanation is another major area. This year we completed some difficult programming that allows us to examine a history in detail of everything NEOMYCIN did when solving a problem. With this foundation, we can now go back and summarize lines of reasoning for any point during the previous consultation. In our first program, completed in 1984, we "translated" steps (metarules) using text strings built into the program. Now we seek to generate these strings automatically by having the program read the metarules and select statements to mention. This project makes significant contributions to text generation research, a somewhat ignored area of natural language research.

C.1.4 The ODYSSEUS Modeling Program

Our third tutorial-related project involves continued development of a modeling program, ODYSSEUS. The purpose of ODYSSEUS is to discover domain knowledge

discrepancies between an application domain knowledge base (e.g. the Neomycin medical knowledge base) and a student or expert problem solver. IMAGE, an earlier modeling program developed in 1982, did not address this problem. The input to ODYSSEUS is the problem solver's patient data requests. When ODYSSEUS watches a student it functions as a student modeling program for GUIDON2 and when it watches an expert it functions as a knowledge acquisition program for HERACLES.

The approach used by ODYSSEUS to detect domain-level discrepancies may be characterized as failure-driven learning by completing explanations. An explanation failure occurs when ODYSSEUS is unable to create a proof tree consisting of instantiated metarules that links an observable student action to a high-level task goal. In creating these proof trees, a top-down simulation first produces a set of plausible high-level goals and updates problem solving state information; then this information is used by a constrained bottom-up generation from the observable action to these high level goals. A explanation failures occurs when no proof tree can be generated for an action and this suggests a domain level discrepancy.

ODYSSEUS resolves this failures in two steps. First, the constraints on proof tree generation are relaxed; this identifies relations in metarules that might be the source of the discrepancy and produces a set of instantiations for each of these relation that are the candidate domain-level discrepancies. Second, a confirmation theory tests these candidate discrepancies for plausibility.

During the last year, ODYSSEUS has been enhanced to operate directly off an arbitrary set of Heracles control metarules; previously the modeling program incorporated knowledge about the particular metarules that were used in Neomycin. This increases the generality and applicability of the program at the cost of a large increase in the search space. Initial validation tests of Odysseus have been conducted and this has revealed that following the strategic reasoning of human problem solvers is crucially dependent on having a very good domain knowledge base. Besides these tests on human problem solvers, a validation methodology called the synthetic agent method has been designed that allows determination of an upper performance bound. During the next year, ODYSSEUS will be completed, integrated with all parts of Guidon including the explanation and Guidon-Manage program, and validated. A case library for the Neomycin domain will be constructed since this plays a crucial role in validation and assessment of the ODYSSEUS approach.

C.1.5 The HERACLES Expert System Shell

The final major effort involves generalizing our expert system tool, HERACLES, so that it can be made available to other research groups who wish to develop knowledge bases which can be tutored by GUIDON2. This project involves a great deal of basic systems programming, including partitioning of files and regrouping of general and knowledge-base-specific constructs. By year end, we were ready to reconfigure a second program built in HERACLES during 1985, called CASTER, to test out the system-building tools developed to date.

A host of smaller projects included:

- Maintenance of our patient library and records of proper program performance.
- Development of a graphics editor for modifying the knowledge base by "correcting" the program's diagnosis of a particular case.
- Development of menu-based knowledge-base retrieval capability. This program constructs menus that bring together details related to some fact the user has just asked a question about.

- More consistent storage and convenient access to "normal values" for patient tests and findings.
- Development of a package for creating, editing, and replaying "scripts" which take the viewer on a tour of some aspect of NEOMYCIN. Useful for documentation, simple lectures, and automatic demonstrations of the program.

C.1.6 Model of Learning

Finally, in a paper described below, we developed a theory of learning by debugging using knowledge of diagnostic strategy and organization of disease knowledge. This theory now forms the foundation for design of GUIDON2. In our current work, we are focusing on the modeling, explanation, and knowledge acquisition capabilities that will allow the tutor to articulate how a diagnostic solution is flawed and how it can be improved using specific domain knowledge. Thus, we are teaching the constraints a good solution must respect, plus giving the students a language for articulating what medical facts are relevant to the case at hand.

C.1.7 Dissemination of results

There were many conferences relating to our work this year. Most notable were the "Tutoring system workshop" in Windermere, England (travel support from the AAAI) and the "Knowledge acquisition workshop" in Banff, British Columbia. Other useful workshops concerned "Higher-level tools" and "Knowledge compilation." Clancey presented prominent papers at each of these workshops and helped organize the middle two. Clancey also presented Guidon/Neomycin work at additional conferences in Milan, London, New Mexico, Arizona, and Florida.

The Macy Foundation Symposium on Cognitive Science and Medical Education in Montreal, run by John Bruer, was extremely valuable for the grantees. Researchers working on medical instruction included: Feltovich, Evans, Hammond, Elstein, and Patel. Small meetings are unusual in this field (AAAI has more than 5000 attendees); the discussions were detailed and illuminating.

Guidon/Neomycin work will be represented in 1987 at Clancey's tutorial on "Evaluating expert system tools" and his tutorial on tutoring systems at IJCAI in Milan.

C.2 Research in Progress

The following projects are active as of May 1987 (see also near-term plans listed in Section III.A):

1. Developing additional instructional programs based on NEOMYCIN;
2. Studying learning in the setting of debugging a knowledge base;
3. Re-implementing the explanation program to use the logic-encoding of the metarules (stating this program in the same task/metarule language so that it might reason about its own explanations);
4. Developing new graphic methods for making presentations from the knowledge base, including tour-like lectures and "dynamic menus" which bring together items relevant to previous user inquiries;
5. Applying the student modeling program, ODYSSEUS, to knowledge acquisition; and
6. Preparing HERACLES, the generalization of NEOMYCIN, for use by other people.

D. Publications Since January 1986

1. Clancey, W.J., Richer, M., Wilkins, D.C., Barnhouse, S., Kapsner, C., Leserman, D., Macias, J., Merchant, A., and Rodolitz, N.: *Guidon-Debug: The student as knowledge engineer.* KSL Working paper 86-34.

2. Clancey, W.J.: *Qualitative Student Models.* **Annual Review of Computer Science.** Palo Alto: Annual Reviews, Report KSL-86-11, Computer Science Dept., May 1986.

3. Clancey, W.J.: *From GUIDON to NEOMYCIN and HERACLES in twenty short lessons: ONR Final Report, 1979-1985. The AI Magazine,* 7(3):40-60, Conference, 1986.

4. Wilkins, D.C., Clancey, W.J., Buchanan, B.G. An overview of the ODYSSEUS learning apprentice. In **Machine Learning: A Guide to Current Research**, eds. T.M. Mitchell, J.G., Carbonell, and R.S. Michalski. New York, Academic Press, pages 369-373. Also KSL-85-26.

5. Clancey, W.J. Intelligent tutoring systems: A tutorial survey. **International Professorship Series, 1985** Academic Press, Inc., London, in press.

6. Wilkins, D.C., Clancey, W.J., and Buchanan, B.G. On Using and Evaluating Differential Modeling in Intelligent Tutoring and Apprentice Learning Systems. In **Intelligent Tutoring Systems: Lessons Learned**, eds. J. Psotka, D. Massey, and S. Mutter. Lawrence Erlbaum Publishers, in preparation. Also KSL-86-62.

7. Clancey, W.J. The knowledge engineer as student: Metacognitive bases for asking good questions. In **Learning Issues in Intelligent Tutoring Systems**, eds. A. Lesgold and H. Mandl, in preparation. Also KSL 87-12.

8. Clancey, W.J. *Viewing knowledge bases as qualitative models.* KSL Working paper 86-27.

9. Clancey, W.J. *Know-how vs. knowledge representation* (extended abstract). *Proceedings of the Workshop on Knowledge Compilation,* Oregon State Technical report, September 1986, pages 1-2.

10. Wilkins, D.C., Clancey, W.J., and Buchanan, B.G., Knowledge Base Refinement Using Abstract Control Knowledge. January, KSL-87-01.

11. Wilkins, D.C., Buchanan, B.G., and Clancey, W.J., *The Global Credit Assignment Problem and Apprenticeship Learning.* January, KSL-87-04.

12. Clancey, W.J. Review of Winograd and Flores's "Understanding Computers and Cognition": A favorable interpretation. *Artificial Intelligence,* December, 1986.

13. Dietterich, T. G., Flann, N. S. and Wilkins, D. C., Machine Learning at IJCAI-85, in *Machine Learning*, Volume 1, No. 2, 1986, 227-242.

14. Karp, P. D. and Wilkins, D. C., An Analysis of the Deep/Shallow Distinction For Expert Systems, KSL-86-32, April 1986, 18 pp.

15. Wilkins, D. C. and Buchanan, B. G., Debugging Rule Sets When Reasoning Under Uncertainty, in *Proceedings of the Fifth National Conference on Artificial Intelligence*, August 1986, 448-454. Also, extended version, KSL-86-30, 20 pp.

16. Wilkins, D. C.,Knowledge Base Debugging Using Apprenticeship Learning Techniques, in *Proceedings of the Knowledge Acquisition for Knowledge-Based Systems Workshop*, November 1986, 40. 0--40. 14. Also, revised version, KSL-86-63, 20 pp.

17. Wilkins, D. C., Clancey, W. J. and Buchanan, B. J., Knowledge Base Refinement Using Abstract Control Knowledge, to appear in *Knowledge Acquisition for Knowledge Based Systems*, edited by J. Boose and B. Gaines, Academic Press. Also to appear in *International Journal of Man-Machine Studies*. Also KSL-87-01, Dec 1986, 12 pp.

18. Wilkins, D. C., Cognitive Diagnosis of Heuristic Classification Problem Solving, *Third International Conference on Artificial Intelligence and Education*, May 1987, pp 57. Also, KSL-86-71, Dec 1986, 2 pp.

E. Funding Support

Contract Title: "A Family of Intelligent Tutoring Programs for Medical Diagnosis"
Principal Investigator: Bruce G. Buchanan, Prof. Computer Science, Research
Associate Investigator: William J. Clancey, Research Assoc. Computer Science
Agency: Josiah Macy, Jr. Foundation
Term: March 1985 to March 1988
Total award: $503,415 direct costs

Contract Title: "Computer-Based Tutors for Explaining and Managing the Process of Diagnostic Reasoning"
Principal Investigator: Bruce G. Buchanan, Prof. Computer Science, Research
Associate Investigator: William J. Clancey, Research Assoc. Computer Science
Agency: Office of Naval Research
ID number: N00014-85-K-0305
Total award: $510,311 total

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Dissemination via SUMEX

We are frequently asked to demonstrate GUIDON-MANAGE, GUIDON-WATCH, and NEOMYCIN to Stanford visitors or at meetings in this country or abroad. Physicians have generally been enthusiastic about the potential of these programs and what they reveal about current approaches to computer-based medical decision making. We use network e-mail through SUMEX to communicate with other researchers worldwide.

B. Sharing and Interaction with Other SUMEX-AIM Projects

GUIDON/NEOMYCIN retains strong contact with the ONCOCIN project, as both are siblings of the MYCIN parent. These projects share programming expertise and utility routines. In addition, the central SUMEX development group acts as an important clearing house for solving problems and distributing new methods.

C. Critique of Resource Management

The SUMEX resources group has provided exemplary service. We have no complaints or suggestions whatsoever.

III. RESEARCH PLANS

A. Project Goals and Plans

Research over the next year will continue on several fronts, including one or more prototype instructional programs.

1. Use GUIDON-MANAGE by medical students to empirically develop the interface and teaching scenario.

2. Integrate the new explanation program into the GUIDON-MANAGE program in order to provide explanations of the operations of tasks invoked by the student.

3. Develop the GUIDON-DEBUG knowledge acquisition program and incorporate its perspective on diagnosis (operators for manipulating the patient-specific model) in feedback provided within GUIDON-MANAGE.

B. Long-term plans

Plans beyond 1988 are uncertain at this time. We expect to make HERACLES available for routine use by people outside of Stanford and explore non-medical applications to broaden our understanding of diagnosis and heuristic classification problem solving.

C. Requirements for Continued SUMEX Use

SUMEX remains the central communications facility for our project--for communication by e-mail and for preparing publications. Research is done on SUMEX-supported Lisp workstations.

D. Requirements for Additional Computing Resources

Within eighteen months, we believe that we will need to upgrade existing workstations purchased in the past few years to incorporate new memory sizes and faster processors. Our experience with color monitors on IBM PC's indicates that the research world must convert to color to fully exploit the potential of computer graphics, especially for knowledge base browsing and editing. There is some question whether academic labs will be left behind by industrial efforts in this respect. We also find that the existing printers are unreliable and of uneven quality. These must be replaced in the near future, perhaps at a higher cost for durability.

E. Recommendations for Future Community and Resource Development

With the proliferation of machine types and the availability of stand-alone machines such as the Macintosh, it is important that the machine be linked for convenient communication by e-mail and conventions be established for automatically translating old publication files into new standard formats.

IV.A.2. MOLGEN Project

MOLGEN - Applications of Artificial Intelligence to Molecular Biology: Research in Theory Formation, Testing, and Modification

Prof. E. Feigenbaum and Dr. P. Friedland
Department of Computer Science
Stanford University

Prof. Charles Yanofsky
Department of Biology
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The MOLGEN project has focused on research into the applications of symbolic computation and inference to the field of molecular biology. This has taken the specific form of systems which provide assistance to the experimental scientist in various tasks, the most important of which have been the design of complex experiment plans and the analysis of nucleic acid sequences. Our current research concentrates on scientific discovery within the subdomain of regulatory genetics. We desire to explore the methodologies scientists use to modify, extend, and test theories of genetic regulation, and then emulate that process within a computational system.

Theory or model formation is a fundamental part of scientific research. Scientists both use and form such models dynamically. They are used to predict results (and therefore to suggest experiments to test the model) and also to explain experimental results. Models are extended and revised both as a result of logical conclusions from existing premises and as a result of new experimental evidence.

Theory formation is a difficult cognitive task, and one in which there is substantial scope for intelligent computational assistance. Our research is toward building a system which can form theories to explain experimental evidence, can interact with a scientist to help to suggest experiments to discriminate among competing hypotheses, and can then revise and extend the growing model based upon the results of the experiments.

The MOLGEN project has continuing computer science goals of exploring issues of knowledge representation, problem-solving, discovery, and planning within a real and complex domain. The project operates in a framework of collaboration between the Heuristic Programming Project (HPP) in the Computer Science Department and various domain experts in the departments of Biochemistry, Medicine, and Biology. It draws from the experience of several other projects in the HPP which deal with applications of artificial intelligence to medicine, organic chemistry, and engineering.

B. Medical Relevance and Collaboration

The field of molecular biology is nearing the point where the results of current research will have immediate and important application to the pharmaceutical and chemical industries. Already, clinical testing has begun with synthetic interferon and human growth hormone produced by recombinant DNA technology. Governmental reports estimate that there are more than two hundred new and established industrial firms already undertaking product development using these new genetic tools.

The programs being developed in the MOLGEN project have already proven useful and important to a considerable number of molecular biologists. Currently several dozen researchers in various laboratories at Stanford (Prof. Paul Berg's, Prof. Stanley Cohen's, Prof. Laurence Kedes', Prof. Douglas Brutlag's, Prof. Henry Kaplan's, and Prof. Douglas Wallace's) and over four hundred others throughout the country have used MOLGEN programs over the SUMEX-AIM facility. We have exported some of our programs to users outside the range of our computer network (University of Geneva [Switzerland], Imperial Cancer Research Fund [England], and European Molecular Biology Institute [Heidelberg] are examples). The pioneering work on SUMEX has led to the establishment of a separate NIH-supported facility, BIONET, to serve the academic molecular biology research community with MOLGEN-like software. BIONET is now serving many of the computational needs of over two thousand academic molecular biologists in the United States.

More generally, our work in qualitative simulation as applied to molecular biology is also relevant to building models of many other medical and biological systems. For example, one Artificial Intelligence researcher (Kuipers) has been applying these techniques to the domain of renal physiology. Other researchers within the KSL are considering applying these techniques to building models of cardio-pulmonary physiology.

C. Highlights of Research Progress

C.1 Accomplishments

During the past year we have concentrated on the qualitative modeling and simulation aspects of the research. Our view is that a well-formulated, multi-level model of a scientific theory is a necessary first step to automated discovery. In addition, we have worked on knowledge acquisition and graphical display of process information and on the description and understanding of the results of laboratory experiments. We have also prepared an in-depth conceptual reconstruction of the biological research which led to the current detailed understanding of the mechanism of attenuation. The highlights of this work are summarized in several categories below.

C.1.1 Qualitative Modeling and Simulation

Our work in qualitative simulation has been directed towards building a program which embodies a theory of the tryptophan system. We have built one model of the system and we are designing a second model based on the successes and failures of the first.

The first model is organized around a set of twenty important state variables of the tryptophan system which we have identified. In addition, it contains descriptions of the causal interactions between these state variables. The novel properties of this model results from the novel representations used for the state variables and the interactions between them.

Our approach to the representation of the values of state variables results from two observations. First, the amount of information biologists have about the values of different state variables varies widely. Second, different amounts of information about a given variable may be available, and of interest, for different problems. Thus, our representation is designed to capture a variety of types of statements about the value of a variable. For example, we can record quantitative information about a variable (`x = .05`), inequality information (`x > 10`), or relative information (`x = 2*y`).

Just as there is a range in the degree of precision with which we might know the value of a given variable, there is an analogous range within which we might know the causal relationship between two variables. Consider that there does exist some function which describes the interactions among any set of variables in our system. Biologists may not

have been able to determine the exact behavior of this function, and hence cannot describe it exactly. Or, we may know its exact behavior, but it may be so complex that we wish to describe it more simply.

Thus, we require a set of representations which allows us to represent the exact form of a function if we have it, or approximations if we do not have it or it is too complex. Relationships among variables are concepts which are represented with several frames, within which all or only some slots may be filled. Relationships between each pair of interacting variables are represented with frames called *Relations*, which describe a unidirectional causal relationship between two variables. For example, we can record any of:

- the sign of a relationship
- whether it is a monotonic relationship
- what the functional form of the relationship is, e.g., linear, higher polynomial, exponential, or unknown
- the sign of the exponent on the input variable
- one or more quantitative coefficients for the relationship

Using these representations we can thus express precisely that (possibly incomplete) knowledge that biologists have about the trp system. We can then define experimental conditions and ask the simulation system to make predictions as to the degree of expression of the genes in the tryptophan operon. For example, we can ask how much expression occurs when the cell is starved of tryptophan, or when tryptophan is in excess. The simulation system propagates the initial experimental conditions through the model in a cyclic fashion to predict how the expression of the operon varies over time.

C.1.2 Process Description and Graphical Display

A system has been built which generalizes our experience in process description by providing a simplified interface for the domain-independent description and animation of process knowledge. The system allows processes to be broken down into component sub-processes and the causal and time-oriented relationships of the subprocesses to be specified. In addition, objects utilized by the processes can be conveniently described and "drawn" with modes and points of interaction among the objects given by the user. All knowledge about processes and objects is automatically stored in the framework of a KEE knowledge base.

After process and object description, the system automatically animates the process by displaying one of several primitive types of interactions among objects in the proper time order dictated by the process knowledge base. This system has been tested on the tryptophan operon domain and its utility is currently being explored in a medical simulation domain.

C.1.3 A Conceptual Reconstruction of the Discovery of Attenuation

Scientific theory formation is a complicated process. The construction of a computer program to reproduce scientific discoveries is one way to study this process. Another way to study the process is by studying the work of actual scientists.

In the past year we have prepared an in-depth study of the discovery of attenuation by Charles Yanofsky and other researchers. We have studied the biological literature extensively and interviewed many scientists involved in the research in order to reconstruct the different conceptual states of knowledge through which the scientists passed in their understanding of the tryptophan operon. By analyzing these states of knowledge and the transitions between them, we have elucidated a number of the strategies and heuristics which these biologists used to generate and choose between

theories of the tryptophan operon. We have related these strategies to both the ideas of different philosophers of science, and to the diagnostic strategies of the Internist medical expert system.

D. Publications

1. Bach, R., Friedland, P., Brutlag, D., and Kedes, L.: *MAXIMIZE, a DNA sequencing strategy advisor.* Nucleic Acids Res. 10(1):295-304, January, 1982

2. Bach, R., Friedland, P., and Iwasaki, Y.: *Intelligent computational assistance for experiment design.* Nucleic Acids Res. 12(1):11-29, January, 1984.

3. Brutlag, D., Clayton, J., Friedland, P. and Kedes, L.: *SEQ: A nucleotide sequence analysis and recombination system.* Nucleic Acids Res. 10(1):279-294, January, 1982.

4. Clayton, J. and Kedes, L.: *GEL, a DNA sequencing project management system.* Nucleic Acids Res. 10(1):305-321, January, 1982.

5. Feitelson, J. and Stefik, M.J.: *A case study of the reasoning in a genetics experiment.* Heuristic Programming Project Report HPP-77-18 (working paper), May, 1977.

6. Friedland, P.: *Knowledge-based experiment design in molecular genetics.* Proc. Sixth IJCAI, August, 1979, pp. 285-287.

7. Friedland P.: *Knowledge-based experiment design in molecular genetics.* Stanford Computer Science Report STAN-CS-79-760 (Ph.D. thesis), December, 1979.

8. Friedland, P., Kedes, L. and Brutlag D.: *MOLGEN--Applications of symbolic computation and artificial intelligence to molecular biology.* Proc. Battelle Conference on Genetic Engineering, April, 1981.

9. Friedland, P.: *Acquisition of procedural knowledge from domain experts.* Proc. Seventh IJCAI, August, 1981, pp. 856-861.

10. Friedland, P., Kedes, L., Brutlag, D., Iwasaki, Y. and Bach R.: *GENESIS, a knowledge-based genetic engineering simulation system for representation of genetic data and experiment planning.* Nucleic Acids Res. 10(1):323-340, January, 1982.

11. Friedland, P., and Kedes, L.: *Discovering the secrets of DNA.* Communications of the ACM, 28(11):1164-1186, November, 1985, and IEEE/Computer, 18(11):49:69, November, 1985.

12. Friedland, P. and Iwasaki Y.: *The concept and implementation of skeletal plans.* Journal of Automated Reasoning, 1(2): 161-208, 1985.

13. Friedland, P., Armstrong, P., and Kehler, T.: *The role of computers in biotechnology.* BIO\TECHNOLOGY 565-575, September, 1983.

14. Iwasaki, Y. and Friedland, P.: *SPEX: A second-generation experiment design system.* Proc. of Second National Conference on Artificial Intelligence, August, 1982, pp. 341-344.

15. Martin, N., Friedland, P., King, J. and Stefik, M.J.: *Knowledge base management for experiment planning in molecular genetics.* Proc. Fifth IJCAI, August, 1977, pp. 882-887.

16. Meyers, S. and Friedland, P.: *Knowledge-based simulation of regulatory genetics in bacteriophage Lambda.* Nucleic Acids Res. 12(1):1-9, January, 1984.

17. Stefik, M. and Friedland, P.: *Machine inference for molecular genetics: Methods and applications.* Proc. of NCC, June, 1978.

18. Stefik, M.J. and Martin N.: *A review of knowledge based problem solving as a basis for a genetics experiment designing system.* Stanford Computer Science Report STAN-CS-77-596, March, 1977.

19. Stefik, M.: *Inferring DNA structures from segmentation data: A case study.* Artificial Intelligence 11:85-114, December, 1977.

20. Stefik, M.: *An examination of a frame-structured representation system.* Proc. Sixth IJCAI, August, 1979, pp. 844-852.

21. Stefik, M.: *Planning with constraints.* Stanford Computer Science Report STAN-CS-80-784 (Ph.D. thesis), March, 1980.

22. Karp, P., and D. Wilkins: *An Analysis of the Deep/Shallow Distinction for Expert Systems.* Stanford University Knowledge Systems Laboratory Report KSL-86-32, 1986.

23. Karp, P., and P. Friedland: *Coordinating the Use of Qualitative and Quantitative Knowledge in Declarative Device Modeling.* Stanford University Knowledge Systems Laboratory Report KSL-87-09, 1987.

24. Round, A.: *QSOPS: A Workbench Environment for the Qualitative Simulation of Physical Processes.* Stanford University Knowledge Systems Laboratory Report KSL-87-37, 1987.

25. Karp, P., and P. Friedland, *A Conceptual Reconstruction of the Discovery of Attenuation.* In Preparation.

E. Funding Support

The MOLGEN grant, which has supported the bulk of this research, is titled: MOLGEN: Applications of Artificial Intelligence to Molecular Biology: Research in Theory Formation, Testing, and Modification. This NSF Grant number MCS-8310236, expired on 10/31/86. The Principal Investigators were Edward A. Feigenbaum, Professor of Computer Science and Charles Yanofsky, Professor of Biology. Additional support for this research has been provided by the Defense Advanced Research Projects Agency, under contract N00039-86C-0033.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

SUMEX-AIM continues to serve as the nucleus of our computing resources. The facility has not only provided excellent support for our programming efforts but has served as a major communication link among members of the project. Systems available on SUMEX-AIM such as EMACS, MM, Scribe and BULLETIN BOARD have made possible the project's documentation and communication efforts. The interactive environment of the facility is especially important in this type of project development.

We strongly approve of the network-oriented approach to a programming environment into which SUMEX has evolved. The ability to utilize Lisp workstations for intensive computing while still communicate with all of the other SUMEX resources has been very valuable to our work. We currently have a satisfactory mode of operation where essentially all programming takes place on the workstations and most electronic communications, information sharing, and document preparation takes place within the mature TOPS-20 environment. The evolution of SUMEX has alleviated most of our previous problems with resource loading and file space. Our current workstations are not quite fast nor sophisticated enough, but we are encouraged by the progress that has been made.

We have taken advantage of the collective expertise on medically-oriented knowledge-based systems of the other SUMEX-AIM projects. In addition to especially close ties with other projects at Stanford, we have greatly benefited by interaction with other projects at yearly meetings and through exchange of working papers and ideas over the system.

The ability for instant communication with a large number of experts in this field has been a determining factor in the success of the MOLGEN project. It has made possible the near-instantaneous dissemination of MOLGEN systems to a host of experimental users in laboratories across the country. The wide-ranging input from these users has greatly improved the general utility of our project.

We find it very difficult to find fault with any aspect of the SUMEX resource management. It has made it easy for us to expand our user group, to give demonstrations to colleagues and to disseminate software to non-SUMEX users overseas.

III. RESEARCH PLANS

A. Project Goals And Plans

Our current work has the following major goals:

1. We will continue our work in qualitative simulation, modeling, and process description. We will continue testing the existing state-variable-based model of the tryptophan operon. In addition, we will construct a new and more general model of the operon. This model will be centered around the objects within this domain (e.g., enzymes, DNA, repressor proteins) and the interactions between them. The current state-variable model makes assumptions about the presence of different objects and the functions of these objects (e.g., that they contain no mutations) which the new model will make both explicit and allow us to change. Essentially, the new model will allow us to dynamically construct new state-variable models based on the presence of different objects and different interactions between these objects. Changing these assumptions is crucial to the discovery process, which involves the postulation of new classes of objects and new classes of interactions between objects.

2. Build a mechanism for postulating extensions or corrections to the current theory: a constrained theory generator. Our conceptual reconstruction of the discovery of attenuation should be of critical help in both this phase and the phases which follow.

3. Build a mechanism for evaluating alternative theories. This would include rating the theories based on plausibility, selectability, completeness, significance, and so on. We hope the evaluation process produces information useful in discriminating among the possible theories.

4. Test the entire structure on the evolving trp operon regulatory system. Experiment with different initial knowledge bases to see how the discovery process is altered by the availability of new techniques, analogous systems, and so forth.

B. Justification and Requirements for Continued SUMEX Use

The MOLGEN project depends heavily on the SUMEX facility. We have already developed several useful tools on the facility and are continuing research toward applying the methods of artificial intelligence to the field of molecular biology. The community of potential users is growing nearly exponentially as researchers from most of the biomedical-medical fields become interested in the technology of recombinant DNA. We believe the MOLGEN work is already important to this growing community and will continue to be important. The evidence for this is an already large list of pilot exo-MOLGEN users on SUMEX.

We support with great enthusiasm the acquisition of satellite computers for technology transfer and hope that the SUMEX staff continues to develop and support these systems. One of the oft-mentioned problems of artificial intelligence research is exactly the problem of taking prototypical systems and applying them to real problems. SUMEX gives the MOLGEN project a chance to conquer that problem and potentially supply scientific computing resources to a national audience of biomedical-medical research scientists.

IV.A.3. ONCOCIN Project

ONCOCIN Project

Edward H. Shortliffe, M.D., Ph.D.
Departments of Medicine and Computer Science
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The ONCOCIN Project is one of many Stanford research programs devoted to the development of knowledge-based expert systems for application to medicine and the allied sciences. The central issue in this work has been to develop a program that can provide advice similar in quality to that given by human experts, and to ensure that the system is easy to use and acceptable to physicians. The work seeks to improve the interactive process, both for the developer of a knowledge-based system, and for the intended end user. In addition, we have emphasized clinical implementation of the developing tool so that we can ascertain the effectiveness of the program's interactive capabilities when it is used by physicians who are caring for patients and are uninvolved in the computer-based research activity.

B. Medical Relevance and Collaboration

The lessons learned in building prior production rule systems have allowed us to create a large oncology protocol management system much more rapidly than was the case when we started to build MYCIN. We introduced ONCOCIN for use by Stanford oncologists in May 1981. This would not have been possible without the active collaboration of Stanford oncologists who helped with the construction of the knowledge base and also kept project computer scientists aware of the psychological and logistical issues related to the operation of a busy outpatient clinic.

C. Highlights of Research Progress

C.1 Background and Overview of Accomplishments

The ONCOCIN Project is a large interdisciplinary effort that has involved over 35 individuals since the project's inception in July 1979. The work is currently in its eighth year; we summarize here the milestones that have occurred in the research to date:

- *Year 1:* The project began with two programmers (Carli Scott and Miriam Bischoff), a Clinical Specialist (Dr. Bruce Campbell) and students under the direction of Dr. Shortliffe and Dr. Charlotte Jacobs from the Division of Oncology. During the first year of this research (1979-1980), we developed a prototype of the ONCOCIN consultation system, drawing from programs and capabilities developed for the EMYCIN system-building project. During that year, we also undertook a detailed analysis of the day-to-day activities of the Stanford Oncology Clinic in order to determine how to introduce ONCOCIN with minimal disruption of an operation which is already running smoothly. We also spent much of our time in the first year giving careful consideration to the most appropriate mode of interaction with physicians in order to optimize the chances for ONCOCIN to become a useful and accepted tool in this specialized clinical environment.

- *Year 2:* The following year (1980-1981) we completed the development of a special interface program that responds to commands from a customized keypad. We also encoded the rules for one more chemotherapy protocol (oat cell carcinoma of the lung) and updated the Hodgkin's disease protocols when new versions of the documents were released late in 1980; these exercises demonstrated the generality and flexibility of the representation scheme we had devised. Software protocols were developed for achieving communication between the interface program and the reasoning program, and we coordinated the printing routines needed to produce hard copy flow sheets, patient summaries, and encounter sheets. Finally, lines were installed in the Stanford Oncology Day Care Center, and, beginning in May 1981, eight fellows in oncology began using the system three mornings per week for management of their patients enrolled in lymphoma chemotherapy protocols.

- *Year 3:* During our third year (1981-1982) the results of our early experience with physician users guided both our basic and applied work. We designed and began to collect data for three formal studies to evaluate the impact of ONCOCIN in the clinic. This latter task required special software development to generate special flow sheets and to maintain the records needed for the data analysis. Towards the end of 1982 we also began new research into a *critiquing model* for ONCOCIN that involves "hypothesis assessment" rather than formal advice giving. Finally, in 1982 we began to develop a query system to allow system builders as well as end users to examine the growing complex knowledge base of the program.

- *Year 4:* Our fourth year (1982-1983) saw the departure of Carli Scott, a key figure in the initial design and implementation of ONCOCIN, the promotion of Miriam Bischoff to Chief Programmer, and the arrival of Christopher Lane as our second scientific programmer. At this time we began exploring the possibility of running ONCOCIN on a single-user professional workstation and experimented with different options for data-entry using a "mouse" pointing device. Christopher Lane became an expert on the Xerox workstations that we are using. In addition, since ONCOCIN had grown to such a large program with many different facets, we spent much of our fourth year documenting the system. During that year we also modified the clinic system based upon feedback from the physician-users, made some modifications to the rules for Hodgkin's disease based upon changes to the protocols, and completed several evaluation studies.

- *Year 5:* The project's fifth year (1983-1984) was characterized by growth in the size of our staff (three new full-time staff members and a new oncologist joined the group). The increased size resulted from a DRR grant that permitted us to begin a major effort to rewrite ONCOCIN to run on professional workstations. Dr. Robert Carlson, who had been our Clinical Specialist for the previous two years, was replaced by Dr. Joel Bernstein, while Dr. Carlson assumed a position with the nearby Northern California Oncology Group; this appointment permitted him to continue his affiliation both with Stanford and with our research group. In August of 1983, Larry Fagan joined the project to take over the duties of the ONCOCIN Project Director while also becoming the Co-Director of the newly formed Medical Information Sciences Program. Dr. Fagan continues to be in charge of the day-to-day efforts of our research. An additional programmer, Jay Ferguson, joined the group in the fall to assist with the effort required to transfer ONCOCIN from SUMEX to the 1108 workstation. A fourth programmer, Joan Differding, joined the staff to work on our protocol acquisition effort (OPAL).

- *Year 6:* During our sixth year (1984-1985) we further increased the size of our programming staff to help in the major workstation conversion effort. The ONCOCIN and OPAL efforts were greatly facilitated by a successful application for an equipment grant from Xerox Corporation. With a total of 15 Xerox LISP machines now available for our group's research, all full-time programmers have dedicated machines, as do several of the senior graduate students working on the project. Christopher Lane took on full-time responsibility for the integration and maintenance of the group's equipment and associated software. Two of our programming staff moved on to jobs in industry (Bischoff and Ferguson) and three new programmers (David Combs, Cliff Wulfman, and Samson Tu) were hired to fill the void created by their departure and by the reassignment of Christopher Lane.

  In addition to funding from DRR for the workstation conversion effort, we have support from the National Library of Medicine which supports our more basic research activities regarding biomedical knowledge representation, knowledge acquisition, therapy planning, and explanation as it relates to the ONCOCIN task domain. We have continued to study the therapy planning process under support from the NLM. This research is led by Dr. Fagan and has concentrated on how to represent the therapy-planning strategies used to decide treatment for patients who run into serious problems while on protocol-described treatment. The physicians who treat these patients often seek out a consultation with the protocol study chairman. Dr. Branimir Sikic, a faculty member from the Stanford University Department of Medicine, and the Study Chairman for the oat cell protocol, collaborated on this project. Janice Rohn joined the ONCOCIN project as data manager and to assist in the knowledge entry process.

- *Year 7:* The seventh year (1985-86) marked several milestones in our research on workstation-based programming. The OPAL knowledge acquisition system became operational, and several new oncology protocols were entered using this system. David Combs was primarily responsible for creating the operational version of OPAL (based on the initial prototype by Joan Differding Walton). As anticipated, we increased the speed and ease with which protocols can be added to the ONCOCIN knowledge base.

  Based on the protocols entered through OPAL, we began experimental testing of the workstation version of ONCOCIN in the Stanford oncology clinic. Clifford Wulfman developed the user interface (based on an initial prototype designed by Christopher Lane). Samson Tu developed the reasoning component (designed originally by Jay Ferguson). Much of their work is built upon an object-oriented system developed for our group by Christopher Lane. We connected the various parts of the system, and demonstrated that we have the capability to run ONCOCIN with the reasoning program and interface program on different machines in the communication network. The current version of the program is currently run on a single workstation, but future versions may take advantage of the multiple machine option. To increase the speed at which we are able to test protocols entered into ONCOCIN, we developed additional programs to test real and synthetic cases without user interaction; these are then reviewed by our collaborating clinicians.

  We also developed a workstation-based program, OPUS, to help clinicians determine which protocols are appropriate for specific patients. OPUS was designed and implemented by Janice Rohn with the assistance of Christopher Lane. We have been using it in the clinic setting since the end of 1985.